

The effect of timing noise on targeted and narrow-band coherent searches for continuous gravitational waves from pulsars

G. Ashton,^{1,*} D.I. Jones,¹ and R. Prix²

¹*School of Mathematics, University of Southampton, Southampton SO17 1BJ*

²*Max Planck Institut für Gravitationsphysik (Albert Einstein Institut), 30161 Hannover, Germany*

(Dated: March 30, 2015)

Most searches for continuous gravitational-waves from pulsars use Taylor expansions in the phase to model the spin-down of neutron stars. Studies of pulsars demonstrate that their electromagnetic (EM) emissions suffer from *timing noise*, small deviations in the phase from Taylor expansion models. How the mechanism producing EM emission is related to any continuous gravitational-wave (CW) emission is unknown; if they either interact or are locked in phase then the CW will also experience timing noise. Any disparity between the signal and the search template used in matched filtering methods will result in a loss of signal-to-noise ratio (SNR), referred to as ‘mismatch’. In this work we assume the CW suffers a similar level of timing noise to its EM counterpart. We inject and recover fake CW signals, which include timing noise generated from observational data on the Crab pulsar. Measuring the mismatch over durations of order ~ 10 months, the effect is for the most part found to be small. This suggests recent so-called ‘narrow-band’ searches which placed upper limits on the signals from the Crab and Vela pulsars will not be significantly affected. At a fixed observation time, we find the mismatch depends upon the observation epoch. Considering the averaged mismatch as a function of observation time, we find that it increases as a power law with time, and so may become relevant in long baseline searches.

PACS numbers: 04.80.Nn, 97.60.Jd, 04.30.Db

I. INTRODUCTION

Rotating neutron stars capable of supporting non-axisymmetric mass distributions will emit continuous gravitational waves¹ (CWs) due to their time-varying quadrupole moments. These may be detectable by next-generation ground-based detectors. The emitted signals can persist for longer than typical search durations, but are weak in amplitude, making them difficult to detect in the noise of the detector.

To find a signal, CW searches use matched filtering techniques such as the \mathcal{F} -statistic [1] which compare the output of the detector with a template. These techniques are powerful provided that the signal and template remain coherent for the duration of the observation. If the signal can be perfectly matched by a template then the signal to noise ratio, used to quantify the detection likelihood, scales as $\rho^2 \propto T_{\text{obs}}$ (e.g. see [2]). This suggests searching over longer observations increases the chances of making a detection.

The templates must model the monotonic spin-down of the source due to the electromagnetic (EM) and gravitational torque; this is done by Taylor expanding the

phase:

$$\phi(t) = \phi_0 + 2\pi \left(f_0(t - t_0) + \frac{\dot{f}_0}{2!}(t - t_0)^2 \right) + \dots, \quad (1)$$

where t_0 is the reference time at which the pulsar frequency, and spin-down parameters $[\phi_0, f_0, \dot{f}_0, \dots]$ are defined. Note that all times refer to the solar system barycentre and we assume the timing model has already correctly accounted for the dispersion measure, proper motion and other parameters as discussed in Edwards et al. [3]. Pulsar astronomers fit this model to observed time of arrivals (TOAs). If the best fit model is accurate enough to track the pulsar to within a single rotation the resulting timing solution is described as *phase-connected*. Often such solutions are capable of tracking the pulsar over durations greater than a year [4]. For gravitational-wave searches, this level of accuracy motivates the use of the same Taylor expansion phase models to account for the spin-down. Pulsar observers measure the frequency f_0^{EM} and higher order coefficients describing the rotation of the pulsar itself. In this work we will consider searches for emission from non-axisymmetric neutron stars at $f_0^{CW} = 2f_0^{EM}$ [5]; from hereon all frequencies and spin-downs refer to the pulsars CW emission.

While Taylor expansion models are on average reliable enough to track the spin-down, pulsars do show deviations. This can either be in the form of glitches, occasional sudden increases in the rotation frequency, or continuous low-frequency variations known as *timing noise*. Glitches and timing-noise may be related phenomena, but are distinguishable by their relative frequency and magnitudes. In this work we restrict our focus to timing noise (we will comment on glitches in sec. IV C)

*E-mail: G.Ashton@soton.ac.uk

¹ Note that in general “continuous waves” can refer to any quasi-monochromatic long-lasting gravitational-wave signals, such as emitted by binaries of white dwarfs, neutron stars or black holes, which would be detectable by LISA or pulsar timing arrays. Here we refer to CWs exclusively in the context of spinning nonaxisymmetric neutron stars as relevant to ground-based detectors.

Timing noise is often represented by structure in the timing residual, which is the difference between the best fit Taylor expansion, typically up to second order in spin-down \dot{f}_0 , and the observed phase. Timing noise refers specifically to deviations from Taylor expansions that are intrinsic to the pulsar and not to systematic errors such as dispersion in the interstellar medium. Hobbs et al. [6] conducted a wide ranging study on timing noise across the pulsar population. They concluded, amongst other things, that timing noise is ubiquitous and inversely correlated to the age of the pulsar. There already exists measures used to quantify the strength of timing noise such as the Δ_8 value introduced by Arzoumanian et al. [7], the generalisation of the Allan variance [8], the covariance function of the residuals [9], and fitting for timing noise as part of the pulsar timing model [10]. These do not convert directly into the effect that timing noise may have on CW searches for pulsars. To quantify this, we need to measure the *mismatch* due to timing noise. This is closely related to the loss of signal to noise ratio due to the imperfect matching between the template and signal (which we define explicitly in section III).

Although a variety of models exist to interpret timing noise [11, 12], there is currently no consensus on a single mechanism. However, for the issue of timing noise and CW searches, we only need to consider the relation between the components of the neutron star which produce the EM and CW signals. This was investigated by Jones [13] who identified three possible scenarios. First, the two signals are strongly coupled: the same timing noise will be observed in both. Second, the two signals are loosely coupled: a similar, but different level of timing noise will be observed in both. Third, timing noise exists only in the EM signal, there is no corresponding variations in the CW signal. Of course these are really three cases from a full spectrum of possibilities which could also include the pulsars CW signal being significantly more noisy than its EM signal.

The significance of timing noise will vary between different types of pulsar CW searches; these can be divided into targeted, narrow-band, directed, and all-sky searches. *Targeted* searches involve a single known pulsar where an estimate of the spin parameters has been obtained from the EM signal. If we assume that the EM and CW signals are strongly coupled, then we can use a *single-template* targeted search. Under this assumption, when the level of timing noise in the EM signal is small, then a single Taylor expansion is sufficient. If instead the level of timing noise is large, then the EM data can be used to account for it; this is done by applying an adapted matched-filtering phase-model that closely follows the observed EM phase model [14]. If instead we assume that the EM and CW signals are loosely coupled, then we should perform a *narrow-band* search in a small area of parameter space. These narrow-band searches aim to allow for small frequency offsets between the EM and CW signals, such as could be caused by free precession, or a finite coupling time between the two compo-

nents of the neutron star [15]. *Directed* searches look for non-pulsing neutron stars predicted by other means such as at the centre of the super-nova remnant Cassiopeia A. An *all-sky* search involves searching over the entire sky for unknown pulsars. For both directed and all-sky searches the lack of EM data necessitates wide bands in the frequency and its derivatives. For fully coherent matched filtering methods these searches can rapidly become computationally prohibitive. To circumvent this, semi-coherent search techniques are used that incoherently combine short fully-coherent sections of data [16]; these will be less sensitive to timing noise. Nevertheless, semi-coherent searches ultimately need to be followed up by targeted fully coherent searches, for which timing noise may be an issue.

For the properties of the CW signal, the most general case is that it will exhibit some timing noise, but it could be different to the timing noise observed in the EM signal. Until a detection is made, we can only make assumptions about how the two are correlated. To probe these assumptions, we will define two special cases corresponding to different sorts of errors in a CW search:

- **Special Case 1:** Timing noise, exactly like that in the EM signal, exists in the CW signal but is not included in the template. This will result in a loss of signal to noise ratio for searches which assumed that timing noise was negligible. The error potentially affects the narrow-band, directed, and all-sky searches since the level of timing noise is unknown. The single template targeted searches will not be effected since they either check that the level of timing noise is negligible, or correct for it using an adaptive phase model.
- **Special Case 2:** Timing noise is included in the template but does not exist in the signal. This will result in a loss of signal to noise ratio for single-template targeted searches that account for timing noise using an adapted phase model (for example Abbott et al. [15]). Instead, these searches will now erroneously introduce timing noise into the template while the signal will be a smooth Taylor expansion.

In this work we will mimic narrow-band and single-template searches to directly simulate special case 1. Specifically, we will inject a fake CW signal which contains a realisation of timing noise, and recover it using templates based on a single global Taylor series. This tests the scenarios in which the timing noise in the CW signal is either exactly coupled to the EM signal, or they are at least similar. However, this also quantifies special case 2 since the signal and template are interchangeable in matched filtering methods. That is, timing noise in the signal but not in the template is equivalent to timing noise in the template but not the signal.

While all known pulsars are potential CW sources, young pulsars are the most promising due to their large

spin-downs (see Abbott et al. [17] for a review). However, it was found by Hobbs et al. [6] that the amount of timing noise is correlated with the spin-down magnitude. This motivated us to study the effect of timing noise on CW searches for neutron stars with large spin-downs.

The realisation of timing noise we will use to investigate timing noise in CW pulsar sources is based on the young Crab pulsar. The Crab is a potentially detectable source of gravitational waves due to its high spin-down rate and it has the highest spin-down upper limit compared to the LIGO noise floor [15]. The EM signal from the Crab is well documented (see sec. II) and contains exceptional levels of timing noise: it was estimated by Jones [13] that such levels of timing noise in the CW signal may cause an issue for current searches.

Several targeted searches have already been performed for CWs from the Crab pulsar. A single-template search for CWs from the Crab pulsar was performed on data collected during the LIGO S5 science run [15]. This search used the Crab ephemeris and an adapted phase model to account for timing noise. In addition to this single-template search, a narrow-band search for signals from the Crab was also performed by Abbott et al. [15] on the S5 data. Another narrow-band search for the Crab was carried out using data from the VIRGO VSR4 science run along with a search for the Vela pulsar [18].

The structure of this paper is as follows. In section II we describe the observational data available from the Crab ephemeris and discuss its relation to CW searches. In section III we describe the signal injection and recovery method. Results from this method are presented in section IV: we begin by considering the effect timing noise has on narrow-band searches, then we consider the mismatch on stretches of data for which narrow-band searches have been performed; we further investigate how the mismatch depends upon epoch; and finally examine how the mismatch depends on the duration of observation. We summarise our results in section V.

II. TIMING NOISE AS DESCRIBED BY THE CRAB EPHEMERIS

The monthly Crab ephemeris [19] provides the phase evolution of the EM signal between 1982 and the present and can be found at <http://www.jb.man.ac.uk/pulsar/crab.html>. It is unlike most timing data for pulsars where a timing model consists of the model parameters (position, spin-down, etc.) given at a single reference time. For the Crab ephemeris, each monthly update consists of the frequency and spin-down coefficients along with a reference time coinciding with the TOA of a pulse at the solar system barycentre. The coefficients are calculated by least-squares fitting of a Taylor expansion to the TOAs. The reference time for each month is chosen as the TOA of the pulse closest to the mid-point; this is done to minimise the average phase error of the local Taylor expansion. The period of a month is short

enough such that these coefficients and equation (1) track the rotational phase during the month.

The Crab ephemeris gives a distinct picture of the variations due to timing noise superimposed on the monotonic spin-down. To illustrate how this manifests itself, figure 1 depicts the frequency evolution in two adjacent months. Notice that a discontinuity occurs at the interface between months. Such discontinuities will occur in the spin-down, frequency, and phase; timing noise can then be described by the magnitude of these jumps. From the Crab ephemeris it can be shown that the distribution of jumps in phase, frequency and spin-down appear to follow standard normal distributions. This is consistent with timing noise models consisting of a large number of small unresolved events accumulating over a month (e.g. the models considered by Cordes and Greenstein [12]).

Timing noise is usually depicted by structure in the phase residuals calculated by removing the best fit Taylor expansion to the phase from the real phase evolution. A best fit Taylor expansion consists of a single set of coefficients f_0 , \dot{f}_0 , and \ddot{f}_0 valid over the *entire* observation period. To make this distinct from the *local* Taylor expansions describing the evolution in each month this will be referred to as the *global template*. In figure 1 we see that if the discontinuity is non-zero, then it is impossible for any global Taylor expansion template to exactly match the local templates in both months. The phase residual, and hence timing noise, results from the inability to match a single global template to all the local ones.




FIG. 1: Illustration of the jumps between 'local' per-month templates in frequency space, defining the frequency jump Δf . This depiction amplifies the order of magnitude of Δf in order to highlight the timing-noise: for the jumps in the ephemeris, Δf is several orders of magnitude smaller than the change in frequency due to spin-down alone.

In this work we aim to quantify the significance of timing noise in CW searches by generating signals from the Crab ephemeris. This is an empirical description of timing noise and so we make no assumptions on the under-

lying astrophysical model.

III. METHOD

We now describe the method to quantify the effect of timing noise on CW searches for signals from isolated pulsar. To be relevant to current pulsar CW search methods, we will base our method on the narrow-band searches of Abbott et al. [15] and Aasi et al. [18]. The results can be interpreted as measuring the consequence of special case 1 on narrow-band and single-template searches; that is we assume the CW signal has a similar level of timing noise as the EM signal and search using global Taylor expansion templates. For this study, a single-template search refers to a single Taylor expansion template and not the adapted phase model proposed by Pitkin and Woan [14].

We begin by generating a CW signal emulating timing noise using the Crab ephemeris. This is done by stringing together month-long smooth Taylor expansion signals. Each month uses the corresponding month from the Crab ephemeris for the Taylor expansion coefficients. The ‘jumps’ at the interface between months constitutes the timing noise. In more detail:

1. From the ephemeris select a period of data consisting of the reference times t_i , frequency f_i , and spin-down \dot{f}_i for each month i .
2. Generate the phase as a function of time from the data and then fit a global Taylor expansion up to \dot{f} for the whole observation time. The fit results in set of interpolated coefficients $[f_0, \dot{f}_0, \ddot{f}_0]$ at a global reference time halfway through the data. These coefficients are used to centre the narrow-band search parameters.
3. We supplement the local monthly data $\{t_i, f_i, \dot{f}_i\}$ with the fixed value of \ddot{f}_0 calculated in the previous step. The phase of the CW signal is always zero at each monthly reference time t_i of the ephemeris, which by construction coincides with a pulse arrival time.

In this process we have assumed that a fixed value of \ddot{f}_0 is sufficient. This can be justified by considering the next term in the Taylor expansion (1) and typical values of $\ddot{f} \sim 10^{-30} \text{ Hz/s}^3$. Over typical search durations $\sim 1 \text{ yr}$ this term contributes less than a radian to the phase, and it can therefore be safely neglected.

We use the **LALSuite** [20] gravitational-wave analysis routines to generate a fake CW signal; for these experiments we work without any simulated detector noise. The standard tool to generate fake CW signals uses single Taylor expansion models. Therefore, to include timing noise in the signal we do the following.

4. We inject each month-long Taylor expansion generated from the Crab ephemeris lasting for only the

duration of that month. This method creates a fake CW signal, lasting several months, which includes timing noise corresponding to the monthly ephemeris.

Once we have produced data, we then use **LALSuite** tools to recover the signal using the \mathcal{F} -statistic [1]. This is a matched filtering method in which the output of the detector is compared to a signal template (see Prix [2] for more details).

Two types of searches are performed: a single template search at the interpolated coefficients $[f_0, \dot{f}_0, \ddot{f}_0]$ and a narrow-band search in f and \dot{f} centred on the interpolated coefficients. These searches were found to be sufficient to find the signal to within a reasonable mismatch, so more sophisticated methods were not required.

The narrow-band consists of a grid of points in f and \dot{f} . As found by Abbott et al. [15] we find searching over \ddot{f}_0 to be unnecessary for this experiment and so it is kept fixed using the value found in step 2 above. The grid spacing is parameterised by m , the one-dimensional maximal mismatch between two adjacent Taylor expansion templates. From Aasi et al. [21] the corresponding grid spacing is given by

$$df = \frac{\sqrt{12m}}{\pi T_{\text{obs}}}, \quad d\dot{f} = \frac{\sqrt{720m}}{\pi T_{\text{obs}}^2}, \quad (2)$$

where T_{obs} labels the observation time.

For the single-template and at each grid point in the narrow-band search, we measure the squared SNR value ρ^2 . In order to quantify the relative loss compared to the perfectly phase-matched squared SNR ρ_s^2 , we define the mismatch in the usual way (e.g. see Prix [22]) as

$$\mu = \frac{\rho_s^2 - \rho^2}{\rho_s^2}. \quad (3)$$

It is well known (e.g. see Prix [2]) that the SNR for a perfectly phase-matched signal is independent of the signal phase evolution. Therefore, in the absence of timing noise the measured value of ρ^2 can reach the maximum value of ρ_s^2 , and the mismatch therefore vanishes in that template. In the presence of timing noise, even the best-matching template will suffer some mismatch, and this effect will increase with the level of timing noise.

In the single-template search, we measure a single mismatch value. The single-template search can also be interpreted to quantify the error made in special case 2, when the template is adapted to account for EM timing noise but none exists in the CW signal. We can think of the narrow-band search as repeating the single-template search over a grid of points; this allows us two degrees of freedom, corresponding to the frequency and spin-down parameters, over which to minimise the mismatch. The grid point with the minimum mismatch, which we denote by μ_{min} , is the best candidate and will be used to quantify the success of the search. Because the narrow-band can minimise the mismatch, μ_{min} must always be equal or smaller than the mismatch in the single-template search.

IV. RESULTS

A. The effect of timing noise on narrow-band searches

We begin by describing how timing noise degrades a narrow-band search. This is done by comparing the result for a signal containing no timing noise with a signal generated from the Crab ephemeris between MJD 45150 and 56668. This period holds no special significance and is used simply to demonstrate the essential features of a signal containing timing noise.


In figure 2 we show the mismatch as a function of parameter space offset for (a) a signal without timing noise, and (b) a signal containing timing noise. The signal without timing noise is injected at the interpolated coefficients $[f_0, \dot{f}_0, \ddot{f}_0]$. Therefore, we find the minimum mismatch with $\mu_{\min} = 0$ at exactly the centre of the grid and the iso-mismatch contours in the local neighbourhood around the origin are well described by ellipses (e.g. see Prix [22]). For the signal with timing noise (b) we notice two distinctive effects: the minimum achievable mismatch μ_{\min} is non-zero, and the iso-mismatch contours around the best-match template are more irregular and less well described by ellipses. In the following we will quantify the effect of timing noise by considering only the location and value of the minimum mismatch grid point in the narrow-band search.

B. Results relevant to recent narrow-band searches


First we consider two particular periods of the Crab ephemeris corresponding to recent narrow-band searches for the Crab: the LIGO S5 period [15] and the VIRGO VSR4 period [18]. The mismatch in the single-template and the minimum mismatch for the narrow-band searches during both periods are listed in table I. For these periods timing noise is found to produce a mismatch of $\approx 1\%$. As expected, the narrow-band mismatch is smaller than the single-template search. The fractional difference between the two searches is relatively small.

Provided that the timing noise observed in the CW signal is at the same level (or less) as that observed in the EM signal, this result signifies that the recent LIGO and VIRGO narrow-band searches would not suffer significantly from the effects of timing noise.

In addition to producing a mismatch, timing noise may result in the best candidate being found at some distance from the centre of the narrow-band search. However, we find that the distance from the centre of the grid is small when compared to the grid spacing used in actual narrow-band searches such as the S5 and VSR4. For the S5 period narrow-band search, we find that the minimum mismatch was a fraction ~ 0.01 of the grid spacing used in the Abbott et al. [15] search. At the resolutions used in real narrow-band searches, the effects of timing noise on the location of the minimum mismatch will not be



(a) Signal without timing noise



(b) Signal with timing noise

FIG. 2: In figure (a) we show the mismatch as a function of parameter space for a signal without timing noise. The injected signal has parameters $[f_0, \dot{f}_0, \ddot{f}_0]$, as a result the mismatch has a minimum at this point.

This can be compared with figure (b) showing the mismatch from a signal including timing noise. The signal is generated from the Crab ephemeris between MJD 45150 and 45668.

evident.

Figure 3 shows the convergence of the measured best mismatch μ_{\min} for the narrow-band search over the S5 period with the value of m . This demonstrates that the non-zero values of μ_{\min} given in table I are not the result of grid coarseness. For signals without timing noise, the measured best mismatch μ_{\min} will have a minimum of $\sim m$ when the putative signal is located halfway between grid points. In the limit of $m \rightarrow 0$ we then expect the measured mismatch to tend to zero. Instead, for a

	Dates MJD	Single template μ	Narrow band μ_{\min}
S5	53673 - 53977	0.00968	0.00933
VSR4	55681 - 55839	0.00659	0.00584

TABLE I: Measurements of the mismatch during the S5 and VSR4 narrow-band search periods.

signal with timing noise we observe a plateau after some initial reduction. This indicates that the grid is now *fully* resolving the variations due to timing noise.




FIG. 3: Measured best mismatch μ_{\min} as a function of grid spacing parameter m (see eqn. (2)), for the Crab pulsar over the S5 period. This demonstrates that μ_{\min} plateaus at a nonzero mismatch suggesting we are resolving a mismatch due to timing noise instead of the effect of finite grid resolution.

C. Minimum mismatch as a function of the observation epoch

We will now investigate how the best mismatch μ_{\min} varies as a function of the observation epoch. We only show the narrow-band search, as the results were found to be very similar for the single-template search. The method consists of measuring the mismatch μ_{\min} in a 6-month window, which is shifted in 1 month intervals over all the available ephemeris data. The observation time of 6 months is chosen to be similar to typical CW search durations. We are restricted to multiples of 1 month by the frequency of updates to the Crab ephemeris.

Timing noise is not the only variability in the spin-down of pulsars - they can also undergo sudden increases in rotation frequency known as *glitches*. The Crab frequently glitches and these are catalogued by Espinoza et al. [23] and available at <http://www.jb.man.ac.uk/pulsar/glitches.html>.

The mechanism which causes a glitch is not well understood and may involve unpredictable variations in the CW signal. As a result, targeted CW searches either avoid periods with known glitches [15], or allow for an arbitrary jump in gravitational wave phase at the time of the glitch [17]. For this work, we ignore the complicating factor introduced by glitches and consider only the effect of timing noise. We do this by omitting windows which include glitches from the search by using the aforementioned glitch catalogue.

We begin by searching in a small 40×40 grid in frequency and spin-down, with a fixed grid space mismatch of $m = 1 \times 10^{-5}$, and the grid spacing as defined in equation (2). It is possible that the minimum mismatch is found at the edge of the narrow-band grid; such candidates are not true local minima in the mismatch. If this is the case, the search is repeated with an increasingly larger grid size, but the same fixed grid spacing. This process continues until we find a minimum mismatch which is not at the edge of the grid.

Figure 4 shows the measured minimum mismatch in the narrow-band search for a sliding 6-month window at the centre of the observation time. The mismatch due to timing noise is the low level noise occurring in between glitches. Greater mismatches are observed in the post-glitch periods; this is expected as the relaxation time after glitches for the Crab is of the order 1 month [4]. We note the presence of an anomalous period of large mismatch for all windows that include the ephemeris time MJD 55362. The cause for this is unclear from the available data, but it may be caused either by a measurement error or a small undetected glitch. In general, we find that the level of mismatch due to timing noise is between $\mu_{\min} \sim 10^{-3} - 10^{-2}$ for these 6-month searches.

D. Averaged minimum mismatch as a function of the observation duration

We can study the averaged behaviour of the mismatch μ_{\min} as a function of time by varying the size of the sliding window in the previous section. This was done for both the narrow-band and single-template searches; the mismatch from the narrow-band search was found to be a fraction $\lesssim 0.1$ smaller on average than the single-template search. We therefore will only present results from the narrow-band search. The shortest possible window ~ 6 months is restricted by the number of points needed to generate a fit to the phase. Setting the upper limit at ~ 17 months retains a statistically meaningful number of points to average over. Having obtained the data from all sliding window sizes in this range we want to analyse the average behaviour as a function of the observation time. Before doing this we filter results in the following ways:

- We do not consider any windows that include or are bounded by glitch events




FIG. 4: Minimum mismatch μ_{\min} found in 6-month sliding window searches as a function of epoch at the centre of window. Vertical dashed lines indicate glitch events as described by Espinoza et al. [23]. The solid vertical line indicates the date MJD 55362, a period of anomalously large mismatch.

- Windows including the anomalous epoch MJD 55362 are omitted. We wish to study the fluctuations due to timing noise, and this period is either an unidentified glitch, or another highly unusual and unrepresentative form of timing noise
- While each entry of the ephemeris is on average valid over a whole month, some months were truncated due to glitches. The sliding window, which works on a fixed number of entries of the ephemeris will occasionally be shorter than average. To ensure we are averaging over windows of a similar length we omit windows for which the observation time differs from the average by 2 weeks.

In figure 5 we plot the averaged minimum mismatch $\langle\mu_{\min}\rangle$ as a function of observation time. This indicates a growth of $\langle\mu_{\min}\rangle$ with observation time resembling a power law.

To quantify the growth of the mismatch, we perform a least-squares fitting to a power law. Fitting the expression:

$$\langle\mu_{\min}\rangle_{\text{fit}} = \kappa \left(\frac{T_{\text{obs}}}{1 \text{ sec}} \right)^n, \quad (4)$$

we find the best fit parameters

$$\kappa = 1.5 \pm 0.8 \times 10^{-23} \quad (5)$$

$$n = 2.88 \pm 0.030. \quad (6)$$

For perfectly matched signals the squared SNR increases linearly [2] with observation time. This suggests




FIG. 5: Averaging the mismatch for sliding window searches and varying the observation times. The points give the mean while the bars correspond to one standard deviation.

that longer observation times yield a greater likelihood of detection. The power law fit with $n > 1$ implies that the average mismatch from Crab timing noise grows faster than the squared SNR. Gains in SNR from longer observation time will therefore eventually be outweighed by the increasing mismatch from timing noise. To estimate when this may occur, we can rearrange equation (3) to give

$$\rho^2 = \rho_s^2 (1 - \mu_{\min}). \quad (7)$$

Substituting the time dependencies for the perfectly matched SNR and the averaged mismatch we have

$$\rho^2 \propto T_{\text{obs}} - \kappa T_{\text{obs}}^{n+1}. \quad (8)$$

Differentiating and solving for T_{obs} yields an expression for the observation time (in seconds) beyond which the ρ^2 value of a signal containing timing noise starts to decrease

$$T_{\text{obs}} = \left(\frac{1}{\kappa(1+n)} \right)^{1/n}. \quad (9)$$

For the fit values from equation (6), this yields a critical observation time of $T_{\text{obs}} \approx 600$ days after which the mismatch exceeds $\langle\mu_{\min}\rangle \approx 0.25$. In this case it is no longer true that further increases in observation time will yield greater detectability.

Jones [13] estimated the maximum time the signal and template would remain coherent given a random walk in frequency. A crude method used a phase residual of 1 rad for the decoherence criteria. For the Crab, this estimates the decoherence time at 200 days. We can improve upon this result by setting a mismatch of 0.1 as the decoherence criteria; using the fit to the averaged mismatch this gives us a decoherence time of $T_{\text{obs}} \approx 400$ days.

The growth of mismatch as a power law is suggestive of random walk timing noise models (see Cordes and Greenstein [12]) for which the rms phase residual also grows as a power law. However, such a scaling in the phase residual must first be converted into a mismatch, which will depend on the search method, before a comparison can be made. In future work we will present a method to achieve this.

V. CONCLUSIONS

We have used observational data on the Crab pulsar to characterise the possible effects of timing noise on coherent targeted single-template and narrow-band continuous gravitational-wave searches for pulsars. This was done by generating fake signals based on the Crab ephemeris data and searching for them using templates without timing noise. Our analysis clarifies the impact for current searches; accordingly, our methods mimic those used by Abbott et al. [15] and Aasi et al. [18].

Our primary results is summarised by Fig. 5: when considering the average mismatch as a function of observation time, we find that the averaged mismatch grows as a power law. In addition to this, we found two interesting aspects when considering the data without averaging over the epoch:

Firstly, for the S5 and VSR4 narrow-band searches, if the timing noise in the CW signal from the Crab is at a similar level (or lower) to that in the EM signal, then we find it will only have a small ($\approx 1\%$) effect on the measured squared SNR of the putative signal. We found the mismatch in single-template searches to be only fractionally larger than the narrow-band searches. This also suggests phase-adapted searches would not be significantly effected if the signal does not contain timing noise.

Secondly, searching over all available Crab data with a 6-month window, we looked at the mismatch as a function of observation epoch. Post glitch periods tend to

admit significant levels of mismatch; this is expected due to the exponential recovery from the glitch. (We also discovered a period around MJD 55362 which has a large mismatch and is not connected to a known glitch). The narrow-band and single-template searches performed similarly in this and subsequent tests. Typically the mismatch due to timing noise for 6-month searches was found to be between 10^{-3} and 10^{-2} .

The scope of this work can be extended to directed and all-sky searches, which target young rapidly spinning down stars which may emit the strongest CWs. These stars are also known to exhibit the highest levels of timing noise and glitch frequently. Crucially the lack of EM data means we cannot be certain a glitch does not occur during the observation and we cannot account for timing noise in the signal. In future work we would like to quantify both these effects and estimate safe upper limits for the search durations. It would also be interesting to consider CW searches for low-mass X-ray binary systems. These are believed to exhibit a stronger form of timing noise known as “spin-wandering”, which constrains the maximal coherence time to the order of a few days before it would lead to a complete loss of SNR, thereby limiting the best achievable sensitivity [24–26].

VI. ACKNOWLEDGEMENTS

GA acknowledges financial support from the University of Southampton and the Albert Einstein Institute (Hannover). DIJ acknowledge support from STFC via grant number ST/H002359/1, and also travel support from NewCompStar (a COST-funded Research Networking Programme). All authors are grateful for useful feedback from members of the Continuous Waves group of the LIGO Scientific Collaboration and the Virgo Scientific Collaboration. This paper has been assigned document number LIGO-P1400207-v1+.

-
- [1] P. Jaranowski, A. Królak, and B. F. Schutz, *Phys. Rev. D* **58**, 063001 (1998), gr-qc/9804014.
 - [2] R. Prix, in *Astrophysics and Space Science Library*, edited by W. Becker (2009), vol. 357 of *Astrophysics and Space Science Library*, p. 651, <https://dcc.ligo.org/LIGO-P060039/public>.
 - [3] R. T. Edwards, G. B. Hobbs, and R. N. Manchester, *MNRAS* **372**, 1549 (2006), astro-ph/0607664.
 - [4] A. Lyne and F. Graham-Smith, *Pulsar Astronomy* (Cambridge University Press, 2012).
 - [5] S. L. Shapiro and S. A. Teukolsky, *Black Holes, White Dwarfs, and Neutron Stars* (John Wiley & Sons, Inc, 1983).
 - [6] G. Hobbs, a. G. Lyne, and M. Kramer, *MNRAS* **402**, 1027 (2010), ISSN 00358711.
 - [7] Z. Arzoumanian, D. J. Nice, J. H. Taylor, and S. E. Thorsett, *Astrophys. J.* **422**, 671 (1994).
 - [8] D. N. Matsakis, J. H. Taylor, and T. M. Eubanks, *Astronomy & Astrophysics* **326**, 924 (1997).
 - [9] W. Coles, G. Hobbs, D. J. Champion, R. N. Manchester, and J. P. W. Verbiest, *MNRAS* **418**, 561 (2011), 1107.5366.
 - [10] L. Lentati, P. Alexander, M. P. Hobson, F. Feroz, R. van Haasteren, K. J. Lee, and R. M. Shannon, *MNRAS* **437**, 3004 (2014), 1310.2120.
 - [11] A. Lyne, G. Hobbs, M. Kramer, I. Stairs, and B. Stappers, *Science* **329**, 408 (2010), 1006.5184.
 - [12] J. M. Cordes and G. Greenstein, *Astrophys. J.* **245**, 1060 (1981).
 - [13] D. I. Jones, *Phys. Rev. D* **70**, 042002 (2004), gr-qc/0406045.
 - [14] M. Pitkin and G. Woan, *Classical and Quantum Gravity* **21**, 843 (2004), gr-qc/0312057.
 - [15] B. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen,

- G. Allen, R. Amin, S. B. Anderson, W. G. Anderson, M. A. Arain, et al., *Astrophys. J.* **683**, L45 (2008), 0805.4758.
- [16] J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, C. Affeldt, et al., *Phys. Rev. D* **85**, 022001 (2012), 1110.0208.
- [17] B. P. Abbott, R. Abbott, F. Acernese, R. Adhikari, P. Ajith, B. Allen, G. Allen, M. Alshourbagy, R. S. Amin, S. B. Anderson, et al., *Astrophys. J.* **713**, 671 (2010), 0909.3583.
- [18] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, et al., *Phys. Rev. D* **91**, 022004 (2015).
- [19] A. G. Lyne, R. S. Pritchard, and F. Graham-Smith, *MNRAS* **265**, 1003 (1993).
- [20] LIGO Scientific Collaboration, *LALSuite: FreeSoftware (GPL) Tools for Data-Analysis* (2014), <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- [21] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, T. Adams, et al., *Physical Review D* **87**, 042001 (2013), 1207.7176.
- [22] R. Prix, *Phys. Rev. D* **75**, 023004 (2007), gr-qc/0606088.
- [23] C. M. Espinoza, A. G. Lyne, B. W. Stappers, and M. Kramer, *MNRAS* **414**, 1679 (2011), 1102.1743.
- [24] J. Aasi et al. (LIGO Scientific Collaboration, VIRGO Collaboration) (2014), 1412.0605.
- [25] P. Leaci and R. Prix, *ArXiv e-prints* (2015), 1502.00914.
- [26] C. Messenger, H. J. Bulten, S. G. Crowder, et al., in preparation (2014), (LIGO DCC-P1400217).